Coalition Formation and Combinatorial Auctions; Applications to Self-organization and Self-management in Utility Computing

Dan C. Marinescu and Ashkan Paya
Computer Science Division
Department of Electrical Engineering and Computer Science
University of Central Florida, Orlando, FL 32816, USA
Email: [dcm, apaya]@cs.ucf.edu

John P. Morrison
Computer Science Department
University College Cork. Cork, Ireland
Email: j.morrison@cs.ucc.ie

March 24, 2015

arXiv:1406.7487v3 [cs.MA] 22 Mar 2015

Abstract

In this paper we propose a two-stage protocol for resource management in a hierarchically organized cloud. The first stage exploits spatial locality for the formation of coalitions of supply agents; the second stage, a combinatorial auction, is based on a modified proxy-based clock algorithm and has two phases, a clock phase and a proxy phase. The clock phase supports price discovery; in the second phase a proxy conducts multiple rounds of a combinatorial auction for the package of services requested by each client. The protocol strikes a balance between low-cost services for cloud clients and a decent profit for the service providers. We also report the results of an empirical investigation of the combinatorial auction stage of the protocol.

1 Introduction and Motivation

Nowadays large farms of computing and storage servers are assembled to support several cloud delivery models including Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS). In such systems users pay only for computing resources they use, similarly to other utilities such as electricity and water.

Computer clouds raise the question of how far we can push the limits of composability of computing and communication systems, while still being able to support effective policies for resource management and their implementation mechanisms. The software, the glue allowing us to build increasingly more complex systems, consists of more and more layers thus, the challenge of controlling large-scale systems is amplified.

Control theory tells us that accurate state information and a tight feedback loop are the critical elements for effective control of a system. In a hierarchical organization the quality of state information degrades as we move from the bottom to the top; *only local information about the state of a server is by definition accurate.* Moreover, this information is volatile, it must be acted upon promptly because the state changes rapidly. Our recent results [22] confirm that hierarchical control has considerably larger overhead than a simple economic model for cloud resource management. The communication complexity of hierarchical control based on monitoring is more than two orders of magnitude higher and consumes a significant fraction of the available bandwidth at all levels of the interconnection network.

Existing solutions for cloud resource management are neither effective nor scalable and they require detailed models of the system and accurate information about the state of individual servers. The polices for cloud resource management must support: (i) admission control; (ii) capacity allocation; (iii) load balancing; (iv) energy optimization, and (v) quality of service (QoS) [21]. Some of the existing and future challenges for cloud resource management which affect these policies are:

-1. The cloud infrastructure is increasingly more heterogeneous; servers with different configurations of multicore processors, GPUs, FPGAs, and data flow engines are expected to become elements of the cloud computing landscape.

-2. The spectrum of cloud services and cloud applications widens. For example, AWS added new services, including Elastic Cache, and Dynamo DB, offers several types of EC2 (Elastic Cloud Computing) profiles including C3 - compute optimized, R3 - memory optimized;

each instance type provides different sets of computer resources measured by vCPUs (vCPU is a hyper-thread of an Intel Xeon core for M3, C3, R3, HS1, G2, and I2). This process is expected to continue at a faster pace to accommodate the so-called Big Data applications.

-3. Cloud over-provisioning demands high initial costs and leads to a low system utilization; this strategy is not economically sustainable [10]. *Elasticity* allows cloud users to increase or decrease their resource consumption based on their needs; elasticity is now based on *over-provisioning*, assembling pools of resources far larger than required to satisfy the average needs. As a result the average cloud server utilization is in the 18% to 30% range [4, 5]. Power consumption of clouds based on over-provisioning is excessive and has a negative ecological impact [4, 26]. A 2010 survey [6] reports that idle or under utilized servers contribute 11 million tonnes of unnecessary CO_2 emissions each year and that the total yearly cost for the idle servers is $19 billion.

-4. The cloud computing landscape is fragmented. CSPs support different cloud delivery models. and this leads to the vendor lock-in; once becoming familiar and storing her data on one cloud it is very costly for the user to migrate to another CSP.

In a large scale-system tensions between local and global objectives exist. These tensions manifest themselves in questions such as: How to balance the individual cost of autonomous servers with global goals e.g., maximizing the CSP payoff? How to adapt the price for services to the actual demand? How to find an equilibrium between system reconfiguration and continuous system availability? Moreover, cooperation must reflect the particular characteristics of the physical organization. Locality is important; indeed, communication across multiple layers of the networking infrastructure is less desirable as the latency increases and the bandwidth decreases. The hypothesis of our research is that self-organization and self-management could address these challenges and provide effective means for cloud resource management.

Informally, self-organization means synergetic activities of elements when no single element acts as a coordinator and the global patterns of behavior are distributed. Self-management means that individuals can effectively set their own goals, make decisions on how to achieve those goals, plan and schedule their activities independently, and evaluate the progress towards these goals. Self-management can lead to faster and more accurate resource management decisions.

Self-management as a result of auctions eliminates the need for a system model and requires only local thus, more accurate information about the state of individual components. This approach has the potential of optimizing the use of resources and allow Cloud Service Providers (CSPs) to offer services at a lower cost for the consumers [13, 19]. Though the virtues of self-management have long been recognized, there is, to our knowledge, no cloud computing infrastructure, based on self-organizing principles and self-management. This is in itself proof of the difficulties to apply these concepts in practice.

Self-management has to be coupled with some mechanisms for coalition formation allowing autonomous agents, the servers, to act in concert. Autonomous systems have to cooperate to guarantee QoS by distributing and balancing the workload, replicate services to increase reliability, and implement other global system policies. Cooperation means that individual systems have to partially surrender their autonomy.

Self-organization cannot occur instantaneously in an adaptive system. It is critical to give the autonomous cloud platforms interconnected by a hierarchy of networks the time to form coalitions in response to services requests thus, self-management requires an effective reservation system. Reservations are ubiquitous for systems offering services to a large customer population, e.g., airline ticketing, chains of hotels, and so on. Existing clouds, e.g., the Amazon Web Services, offer both reservations and spot access, with spot access rates lower than those for reservations.

The solution discussed in this paper involves concepts, policies, and algorithms from several well-established areas of economics and computer science: self-organization and self-management of complex systems; coalition formation and virtual organizations; auction theory and practice; and system organization and computer architecture. We discuss related work and our contributions in Section 2 and in Section 3 we describe the system model. Algorithms for the formation of sub-coalitions and for clock-proxy auction are the subjects of Sections 4 and 5, respectively. The results of a simulation experiment and the conclusions of our work are presented in Sections 6 and 7.

2 Related Work

The present and future challenges outlined in Section 1 motivate the search for effective and scalable policies and mechanisms for cloud resource management [7, 9, 18, 23, 25, 29, 36]. In this section we survey some of the research in this area focused on market mechanisms.

Coalition formation. Informally, a *coalition* is a group of entities which have agreed to cooperate for achieving a common goal. A *virtual organization* involves entities that require a communication infrastructure and dedicated software to support their activities. *Coalition formation* is a widely used method for increasing the efficiency of resource utilization and for providing convenient means to access these resources [24]. In recent years, the emergence of large-scale electronic markets, grid and cloud computing, sensor networks, and robotics have amplified the interest in coalition formation and virtual organizations [16, 17, 31]. For example,

self-organization of sensor networks through bottom-up coalition formation is discussed in [20, 32].

Different aspects of resource management in computational grids including load balancing, job-allocation, and scheduling, as well as revenue sharing when agents form coalitions or virtual organizations are discussed in [8, 14, 15, 27, 33, 35]. Grid resource allocation is modeled as cooperative games [15] or non-cooperative games [27]. Resource co-allocation is presented in [35].

There is little surprise that the interest in coalition formation migrated in recent years from computational grids to cloud resource management. The vast majority of on-going research in this area is focused on game-theoretic aspects of coalition formation for cloud federations. A *cloud federation* is of a set of CSPs collaborating to provide services to a cloud user community.

A stochastic linear programing game model for coalition formation is presented in [25]; the authors analyze the stability of the coalition formation among cloud service providers and show that resource and revenue sharing are deeply intertwined. An optimal VM provisioning algorithm ensuring profit maximization for CSPs is introduced in [9].

A cloud federation formation described as a hedonic game and focused on the stability and the fairness of the game is discussed in [23]. The profit maximization for each federation is formulated as an integer programming problem (IP) and the game is augmented with a preference relation over the set of federations. The paper assumes that the Virtual Machines (VMs) contributed by each CSP to a federation are characterized by several attributes, $a \in \mathcal{A}$ including the number of cores, the amount of memory and of secondary storage. The IP problem for CSP $\mathcal{C}_i$ in federation $\mathcal{F}$ is formulated as $\max \sum_{\mathcal{C}_i \in \mathcal{F}} \sum_{j=1}^{n} n_{i,j}(p_j - c_{i,j})$ subject to the set of conditions $\sum_{j=1}^{n} q_j^a n_{i,j} \leq A_i,\ \forall a \in \mathcal{A}$ and $\sum_{\mathcal{C}_i \in \mathcal{F}} n_{i,j} = r_j$ with: $n_{i,j}$ - the number of VMs of type j; p_j - the price for a VM running an instance of type j; $c_{i,j}$ - the cost of an instance of type j provided by $\mathcal{C}_i$; q_j^a - the quantity of resource of type a in a VM of type j; A_i - the total amount of resource of type a offered by $\mathcal{C}_i$; and r_j - the number of VMs of type j requested. The paper adopts a payoff division based on the Banzhaf value [23].

A combinatorial coalition formation problem is described in [17]. The paper assumes that a seller has a price schedule for each item. The larger the quantity requested, the lower is the price a buyer has to pay for each item; thus, buyers can take advantage of price discounts by forming coalitions. A similar assumption is adopted by the authors of [16] who investigate systems where the negotiations among deliberate agents are not feasible due to the scale of the system. The paper proposes a macroscopic model and derives a set of differential equations describing the evolution in time of coalitions with a different number of participants. The results show that even a low rate of leaving away participants allows a coalition to achieve a steady state.

An algorithm to find optimal coalition structures in cooperative games by searching through a lattice like the one in Figure 2, was introduced by [30]. A more refined algorithm is described in [28]; in this algorithm the coalition structures are grouped according to the so-called *configurations* reflecting the size of the coalitions.

Auctions. Auctions are a widely used mechanism for resource allocation [11, 12]. Among the numerous applications of auctions are: the auctioning of airport take-off and landing slots, spectrum licensing by the Federal Communication Commission (FCC), and industrial procurement. An online auction mechanism for resource allocation in computer clouds is presented in [35].

A *combinatorial auction* is one where a buyer requires simultaneous access to a *package* of goods. An auction allows the seller to obtain the maximum feasible profit for the auctioned goods; it is organized by an *auctioneer* for every *request* of a consumer. A *proxy* is an intermediary who collects individual bids from the buyers participating at an auction, computes the total cost of the package from the bids, and communicates this price to the auctioneer. A vast literature including [1, 2, 3, 34] covers multiple aspects of combinatorial auctions including bidding incentives, stability, equilibrium, algorithm testing, and algorithm optimality.

Package bidding assumes that a seller offers $\mathcal{N}$ different types of items. A buyer bids for packages of items. A *package* is a vector of integers $\mathcal{Z} = \{z_1, z_2, \ldots, z_{\mathcal{N}}\}$ which indicates the quantity of each item in the package; the price of items is given by $\mathcal{M} = \{m_1, m_2, \ldots, m_{\mathcal{N}}\}$.

Package bidding can be traced back to generalized Vickerey auctions based on the Vickerey-Clarke-Groves mechanisms [11, 12]. In Vickerey auctions a bidder reports its entire demand schedule. The auctioneer then selects the allocation which maximizes the total value of the package and requires a bidder to pay the lowest bid it would have made to win its portion of the final allocation, considering all other bids.

In an *ascending package auction (APA)* there are $\mathcal{K}$ participants identified by an index, $k = 0$ is the seller and $k = 1, 2, \ldots, \mathcal{K}$ are the buyers [1]. Each buyer has a *valuation* vector $v_i = (\nu_i(z), z \in [0, \mathcal{M}])$; $\nu_k(z)$ represents the value of package z to the bidder k. Some of the rules for this type of auction are: all bids are firm, a bid cannot be reduced or withdrawn; the auctioneer identifies after each round the set of the bids that maximize the total price, the so-called *provisional winning bids*. The auction ends when a new round fails to elicit new bids; then the provisional winning bids become the winers of the auction.

In an ascending package auction a bidder can be deterred from bidding for the package she really desires by the threat that competitors could drive prices up; this would threaten the equilibrium. This problem does not exist in *ascending proxy auctions* when each bidder instructs a proxy agent to bid on her behalf [1]. The proxy accepts as input the bidder's valuation profile and bids following a "sincere strategy." Nash equilibrium can be

reached when the bid increments are negligibly small [1].

In a *clock auction* the auctioneer announces prices and the bidders indicate the quantities they wish to buy at the current price. When the demand for an item increases, so does its price until the there is no excess demand. On the other hand, when the offering exceeds the demand, the price decreases [1]. In a clock auction the bidding agents see only aggregate information, the price at a given time, and this eliminates collusive strategies and interactions among bidding agents. The auction is *monotonic*, the amounts auctioned decrease continually and this guarantees that the auction eventually terminates. When the price of a package can be computed as the sum of products of prices and quantities it is said that auction benefits from *linear pricing*.

The *clock-proxy-auction* is a hybrid auction based on an iterative process with two phases [3]. A *clock* phase is followed by a *proxy* round. During the proxy round the bidders report the values they have submitted to the proxy which in turn submit bids for the package to the auctioneer. A bidder has a single opportunity to report the quantity and the price to the proxy, bid withdrawals are not allowed, and the bids are mutually exclusive. The auctioneer then selects the winning bids that maximize the seller's profit.

The contribution of this paper. The reservation system we propose has two stages; coalitions of servers are formed periodically during the first and in the second the coalitions participate in combinatorial auctions organized in each allocation slot. To our knowledge this is the first attempt to address cloud self-organization and resource management based on coalition formation and combinatorial auctions when individual servers learn from past behavior, see Figure 1.

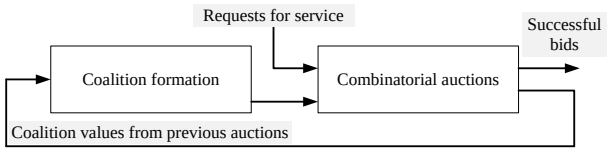


Figure 1: A protocol with two stages; feedback about past values of individual coalitions is used to determine the value of individual coalition structures as shown in Section 4.

We discuss coalition formation for a realistic model of the cloud infrastructure, hierarchical organization, while most of the research reported in the literature is focused on coalition formation for cloud federations. The coalition formation problem has different formulations and different constraints in the two cases.

At this time individual CSPs believe that they have a competitive advantage due to the unique value of their services and are not motivated to disclose relevant information about the inner working of their systems as we have re-discovered when investigating the energy consumption of AWS instances [26]. Thus, the practical realization of cloud federations seems a rather remote possibility [21].

A rare glimpse at the architecture of a cloud is provided in [5] and we are taking advantage of it to base our research on a realistic model of the cloud infrastructure. We investigate coalition formation subject to the physical constraints of the hierarchical cloud organization model. As more diverse applications, including Big Data applications, are likely to use computer clouds, the demand for computing resources allocated to a single application will increase and could be considerably larger than any server can provide; only coalitions of servers will be capable to offer such resources. It is critical for the members of a coalition to communicate effectively; this requires coalition member to be in close proximity of each other in a system consisting of a hierarchy of networks with different bandwidth and latency. This adds additional constrains to the coalition formation protocol.

To respond to the needs of increasingly more complex applications consisting of multiple phases and requiring workflow management, CSPs are already offering workflow management services such as SWS (Simple Workflow Management) and EBS (Elastic Bean Stock) at AWS. Different phases of an application may require coalitions of servers with different types of resources and this is the reason why we decided to investigate combinatorial auctions where packages of items are auctioned.

3 System Model

System architecture. We assume a hierarchical organization of the cloud infrastructure similar to the one described in [5]. A data center consists of multiple warehouse-scale computers (WSCs), each WSC has multiple cells, each cell has multiple racks and each rack houses multiple servers. A WSC connects $50,000$ to $100,000$ servers and uses a hierarchy of networks. The servers are housed in racks; typically, the 48 servers in a rack are connected by a 48 port Gigabit Ethernet switch. The switch has two to eight up-links which go to higher level switches in the network hierarchy [5]. The bandwidth to communicate outside the rack is much smaller than the one within the rack; this has important implications for resource management policies and becomes increasingly difficult to address in systems with a large number of servers.

Model assumptions. For simplicity we assume that the racks are homogeneous, they have identical processors with the same number of cores and an identical configuration of GPUs, FPGAs, workflow engines, or other hardware, the same amount of main storage, cache, and secondary storage. We also assume that all servers in a rack are identically configured and support the same type of services. The same service may be offered by

multiple racks; for example, multiple racks could offer configurations with GPUs.

The system we envision supports a *reservation system* and *spot* resource allocation. The reservation system has two stages: (A) coalition formation, and (B) combinatorial auctions. The spot allocation is done through a bidding process for each type of service. The time is quantified, reservations are made as a result of auctions carried out at the beginning of each *allocation slot* of duration τ; for example, a allocation slot could be one hour.

Coalition formation. The rational for coalition formation is that applications may need resources beyond those provided by an individual server. For example, a Map-Reduce application may require a set of 20 servers during the Map phase to process a data set of several PB (Petabytes). If the algorithm requires the servers to communicate during this phase then the application should start at the same time on all servers and run at the same pace, a condition known as *co-scheduling*. Co-scheduling is only feasible if the set of 20 servers form a coalition dedicated to the application; moreover, the hardware configuration of the coalition members should be optimal for the algorithms used by the application, e.g., have attached GPUs.

Combinatorial auctions. Combinatorial auctions allow cloud users making the reservations to acquire packages consisting of coalitions of servers with different types and amounts of resources. Combinatorial auctions are necessary because different phases of an application may require systems with different configurations or systems supporting different functions. In our previous example the Reduce phase of the Map-Reduce application may require several servers with a very large amount of secondary storage.

4 Coalition Formation

First, we discuss the formulation of the coalition formation problem as a cooperative game. Then we introduce the algorithms for determining the optimal coalition structure and for coalition formation in the context of our model.

Coalition formation as a cooperative game. The coalition formation is modeled as a cooperative game where the goal of all agents is to maximize the reward due to the entire set of agents. We consider a set of N servers $\{s_1, s_2, \ldots, s_N\}$, located in the same rack.

A *coalition* $\mathbb{C}_i$ is a non-empty subset of N. A *coalition structure* is set of m coalitions $\mathbb{S} = \{\mathbb{C}_1, \mathbb{C}_2, \ldots, \mathbb{C}_m\}$ satisfying the following conditions

$$\bigcup_{i=1}^{m} \mid \mathbb{C}_i \mid = N \;\text{ and }\; i \neq j \;\Rightarrow\; \mathbb{C}_i \bigcap \mathbb{C}_j = \emptyset. \quad (1)$$

Figure 2 shows a lattice representation of the coalition structures for a set of four servers s_1, s_2, s_3 and s_4. This lattice has four levels, $L1, L2, L3$ and $L4$ containing the coalition structures with $1, 2, 3$ and 4 coalitions, respectively. In general, the level k of a lattice contains all coalition structures with k coalitions; the number of of coalitions structures at level k for a population of N agents is given by the Sterling Number of Second Kind:

$$\mathcal{S}(N,k) = \frac{1}{k!}\sum_{i=0}^{k}(-1)^i\binom{k}{i}(k-i)^N. \quad (2)$$

In the case illustrated in Figure 2 $N = 4$ and the number of coalition structures at levels $L1 - L4$ are $1, 7, 6, 1$, respectively.[1] The total number of coalition structures with N agents is called the Bell number

$$\mathcal{B}(N) = \sum_{k=0}^{N}\mathcal{S}(N,k) = \sum_{k=0}^{N}\frac{1}{k!}\sum_{i=0}^{k}(-1)^i\binom{k}{i}(k-i)^N. \quad (3)$$

The number of coalitions structures increases exponentially with the number of agents. For example, for $N = 40$, a typical number of servers in a rack, the logarithm of the number of coalition structures is close to 10^{35} and $\mathcal{S}(40, 14) = 3.58598722556218034914285 54E + 34$. The logarithm of number of coalitions is close to $E + 10$.

Searching for the optimal coalition structure $\mathbb{C}$ is computationally challenging due to the size of the search space. The first step for determining the optimal coalition structure is to assign a *value* v reflecting the utility of each coalition. The second step is the actual coalition formation.

Rack-level coalition formation. Recall from Section 3 that in our model a rack is homogeneous, all servers have an identical configuration. This realistic assumption simplifies considerably the complexity of the search for an optimal coalition structure as the servers are indistinguishable from one another.

The second important observation is that we have a system with two stages and feedback, see Figure 1. In the second stage the coalitions created during the first stage are included in successfully auctioned packages thus, we can determine precisely the value of all coalitions structures. The third important observation is that only *available servers*, servers with no commitments for the current slot, can participate to coalition formations and then to the auction organized in that slot; call $N_a \leq N$ the number of available servers.

An elected *rack leader* collects information about all *successful coalitions* - coalitions that have been included in packages auctioned successfully during a window of w successive past allocation slots. The current rack-leader records an entry for the corresponding *partial coalition*

[1] For $N = 5$ and $N = 6$ the Stirling Numbers of the Second Kind are respectively $1, 15, 25, 10, 1$ and $1, 31, 90, 65, 15, 1$.

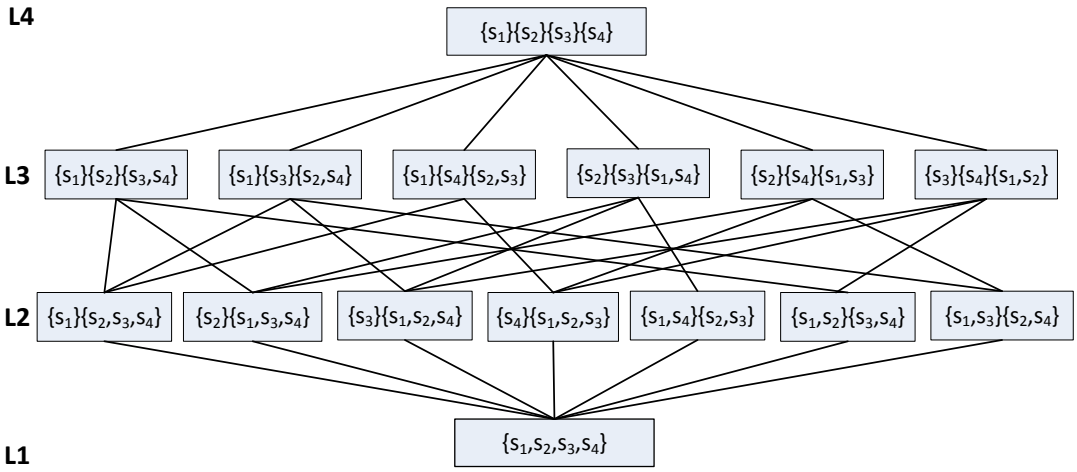


Figure 2: A lattice with four levels $L1, L2, L3$ and $L4$ shows the coalition structures for a set of 4 servers, s_1, s_2, s_3 and s_4. The number of coalitions in a coalition structure at level L_k is equal to k.

structure (PCS) including n_k - the *coalition size*, m_k - the *multiplicity* of occurrence, the value $\bar{v}_k$ calculated as the average price over all auctions when a PCS including a coalition of size n_k was part of a *package* successfully auctioned during the past w allocation slots.

Call $\mathcal{L}$ the *PCL-list*. For a window of size w the list $\mathcal{L}$ is the list of all triplets $\mathcal{L}_k = [n_k, m_k, \bar{v}_k]$ ordered first by $1 \leq n_k \leq N_a$ then by m_k. The list includes only entries $\mathcal{L}_k$ with $\bar{v}_k > 0$. Given N_a a *coalition structure (CS)* $\mathbb{S}_k$ among the entries $\mathcal{L}_{k1}, \mathcal{L}_{k2}, \ldots\ldots, \mathcal{L}_{kn}$ is *feasible* if $\sum_j n_k \times m_k = N_a$. Then the value of the coalition structure $\mathbb{S}_k$ is $v_k = \sum_j \bar{v}_j$. Note that we force the formation of coalitions involving all available servers. An example of a PCS list $\mathcal{L}$ follows

```
-----------------------------------------------
a  [1,4,35]     \* 4 PCS of 1-server {s}
b  [1,15,682]   \ *15 PCS of 1-server {s}
       .........
c  [2,3,78]      \* 3 PCS of  2-servers {s,s}
       ........
d  [3,2,502]    \* 2 PCS of 3-servers {s,s,s}
e  [3,4,812]    \* 4 PCS of 3-servers {s,s,s}
       .........
f  [16,1,751]   \* 1 PCS of 16-servers {s,...s}
g  [16,2,740]   \* 2 PCS of 16-servers {s,...s}
       .........
-----------------------------------------------
```

In this example some of the feasible coalitions structures when $N_a = 16$ are: $\mathbb{S}_g$ with $v_g = 751$; $\mathbb{S}_{a,b}$ with $v_{a,b} = 35 + 682 = 712$; $\mathbb{S}_{a,e}$ with $v_{a,e} = 35 + 812 = 837$; $\mathbb{S}_{a,c,d}$ with $v_{a,c,d} = 35 + 78 + 502 = 615$, and so on. Note that the value of a coalition reflects also the length of time the coalition was active in response to successful auction. We see that a PCS of 15 coalitions of 1 server have been active for larger number of slots than a PCS of 4 coalitions of 1 server. The value attributed to a coalition of k servers is a distributed equally among the servers; the value of a package of several coalitions auctioned successfully is divided among the coalitions based on the resource supplied by each one of them.

Coalition formation. The protocol for coalition formation proceeds as follows:

1. Server s_i sends to the current rack leader:
 (a) A vector $([\nu_i^1, \beta_i^1], [\nu_i^2, \beta_i^2], \ldots\ldots [\nu_i^N, \beta_i^N])$ with $\nu_i^k, 1 \leq k \leq N$ the total value due to the participation of s_i in successful coalitions, of k servers and β_i^k a bit vector with w components with $\beta_i^{k,j} = 1$ if s_i was included in a successful coalition of k servers in slot j of window w.
 (b) Availability, $a_i = 1$ if available, 0 otherwise.
2. After receiving the information from all servers the current rack leader:
 (a) Determines $N_a = \sum_{i=1}^{N} a_i$.
 (b) Computes $m_k = \sum_{i=1}^{N_a} \sum_{j=1}^{w} \beta_i^{k,j}, 1 \leq k \leq N$.
 (c) Computes $\bar{v}_k = \sum \nu_i^k$
 (d) Computes the optimal coalition structure.
 (e) Assigns a server to coalition of size k a based on the values ν_i^k.
 (f) Chooses the best performer as the next coalition leader. The best performer is the one with the largest value $\sum_j \nu_i^j$.

Finding the optimal CS requires at most L operations with L the size of the PCL-list. The system starts with a predetermined coalition structure and coalition values.

5 A Reservation System Based on a Combinatorial Auction Protocol

The protocol introduced in this section targets primarily the *IaaS* cloud delivery model represented by Amazon Web Services (AWS). Reservation systems are currently used by CSPs. For example, AWS supports reservations as well as spot allocation and offers a limited number of instance families, including M3 (general purpose), C3 (compute optimized), R3 (memory optimized), I2 (storage optimized), G2 (GPU) and so on. An instance is a package of system resources; for example, the `c3.8xlarge` instance provides 32 vCPU, 60 GiB of memory, and 2×320 GB of SSD storage. The resources auctioned are supplied by coalitions of servers in different racks and the cloud users request packages of resources.

The combinatorial auction protocol is inspired by the clock-proxy auction [3]. The clock-proxy auction has a clock phase, where the price discovery takes place, and a proxy phase, when bids for packages are entertained. In the original clock-proxy auction there is one seller and multiple buyers who bid for packages of goods.

For example, the airways spectrum in the US is auctioned by the FCC and communication companies bid for licenses. A package consist of multiple licenses; the quantities in these auctions are the bandwidth allocated times the population covered by the license. Individual bidders choose to bid for packages during the proxy phase and pay the prices they committed to during the clock phase.

Our protocol supports auctioning service packages; a packages consist of combinations of services in one or more time slots. The items sold are services advertised by coalitions of autonomous servers and the bidders are the cloud users. Each service is characterized by

1. A *type* describing the resources offered and the conditions for service,

2. The time slots when the service is available.

Protocol specification. The terms used to describe the protocol are discussed next. An *allocation slot* (AS) is a period of fixed duration, e.g., one hour, that can be auctioned. An *auction*, $\mathbb{A}^t$, is organized at time t if there are pending reservation requests which require immediate attention. Figure 3 shows two consecutive auctions at times t and s; during the first slot of auction $\mathbb{A}^t$ new reservation requests are received and the allocation slot AS_2^t is not fully covered; this slot becomes AS_1^s for $\mathbb{A}^s$.

A service $\mathcal{A}$ is described by a relatively small number of *attributes*, $\{a_1, a_2, \ldots, \}$. Each attribute a_i can take a number of distinct values, $v_i = \{v_{i,1}, v_{i,2}, \ldots\}$. The first attribute is the coalition size or equivalently the number of vCPS provided; other attributes could be the type of service and server architecture with two values "32-bit" and "64-bit;" another attribute could be "organization" with values "vN" (von Neumann), "DF" (data-flow), or "vN-GPU" (vN with graphics co-processor).

Call $\mathcal{S}^t$ the set of services the clients want to reserve during auction $\mathbb{A}^t$

$$\mathcal{S}^t = \{S_1^t, S_2^t, \ldots, S_{\nu^t}^t\} \text{ with } S_i^t = [sId, (a_j, v_{j,k})] \quad (4)$$

A *reservation bundle,* $\alpha_{i,j}^t \subset \mathcal{S}^t$, is the set of services requested by client i in slot j of auction $\mathbb{A}^t$

$$\alpha_{i,j}^t = \{(S_{i,j,1}^t, r_{i,j,1}^t), (S_{i,j,2}^t, r_{i,j,2}^t), \ldots\} \quad (5)$$

with $r_{i,j,l}^t$ a measure of the quantity; for example, if the attribute is "service intensity" the quantity is the number of vCPUs.

An *advertised bundle,* $\beta_{k,j}^t \subset \mathcal{S}^t$, is the set of services advertised by coalition k in slot j of auction $\mathbb{A}^t$

$$\beta_{k,j}^t = \{(S_{k,j,1}^t, q_{k,j,1}^t, p_{k,1}^t), (S_{k,j,2}^t, q_{k,j,2}^t, p_{k,2}) \ldots\} \quad (6)$$

with $q_{k,j,l}^t$ a measure of the quantity of service l and $p_{k,l}$ the price per ECU of service S_l^t determined by coalition k. A *package*, $\mathcal{P}_i^t$ is a set of reservations for services requested by client i for slots $j_1, j_2, \ldots$ during auction $\mathbb{A}^t$.

$$\mathcal{P}_i^t = \{\alpha_{i,j_1}^t, \alpha_{i,j_2}^t, \ldots\} \quad (7)$$

The clock phase. Figure 4 illustrates the basic idea of a clock phase: the auctioneer announces prices and the bidders indicate the quantities they wish to buy at the current price. When the demand for an item increases, so does its price until there is no excess demand; on the other hand, when the offering exceeds the demand, the price decreases.

During the clock phase of auction $\mathbb{A}^t$ the price discovery is done for each time slot and for each type of service; a clock runs for each one of the κ^t slots and for each one of the ν^t services. Next we describe the clock phase for service S_l^t in slot j. Assume that there are n coalitions $\mathbb{C} = \{\mathbb{C}_1, \mathbb{C}_2, \ldots, \mathbb{C}_n\}$ offering the service and m requests for reservations from clients $\mathbb{D} = \{\mathbb{D}_1, \mathbb{D}_2, \ldots, \mathbb{D}_m\}$.

A clock auction starts at clock time $t = 0$ and at price per unit of service for S_l

$$p_l^0 = \min_{\mathcal{C}_k}\{p_{k,l}\} \quad (8)$$

Call $\mathcal{C}_0$ the available capacity at this price and $\mathcal{D}_0$ the demand for service S_l^t offered at price p_l^0 in slot j

$$\mathcal{C}_0 = \sum_{k=1}^{n} q_{k,j,l}^t \text{ and } \mathcal{D}_0 = \sum_{i=1}^{m} r_{i,j,l}^t. \quad (9)$$

If $\mathcal{C}_0 < \mathcal{D}_0$ the clock c advances and the next price per unit of service is

$$p_l^1 = p_l^0 + \mathcal{I} \quad (10)$$

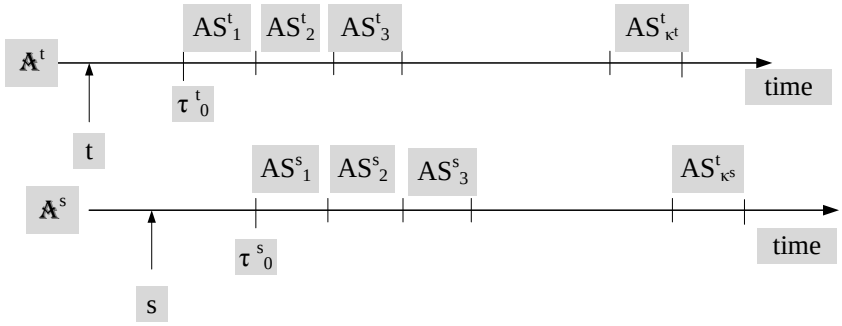


Figure 3: Auctions $\mathbb{A}^t$ and $\mathbb{A}^s$ conducted at times t and s, respectively. τ_0^t and τ_0^s are the start of the first allocation slots, AS_1^t and AS_1^s of the two auctions. The number of slots auctioned in each case are κ^t and κ^s, respectively.

with $\mathcal{I}$ the price increment decided at the beginning of auction. There is an ample discussion in the literature regarding the size of the price increment; if too small, the duration of the clock phase increases, if too large, it introduces incentives for gaming [3].

The process is repeated at the next clock value starting with the new price. The clock phase for service S_l^t and slot j terminates when there is no more demand.

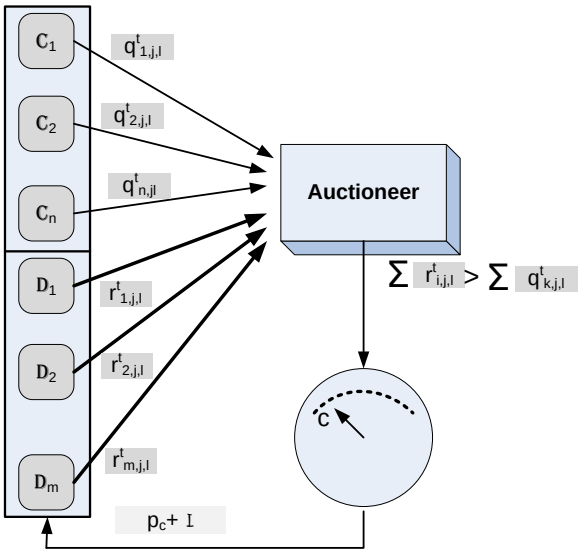


Figure 4: The clock phase for service S_l^t and slot j. The starting price is p_l^0 given by Equation 8. The clock advances and the price increases from p_c to $p_c + \mathcal{I}$ when the available capacity at that price given by Equation 9 is exhausted; the demand is given by Equation 9.

The proxy phase. In a traditional clock-proxy auction the bidders do not bid directly, they report to a proxy the price and the quantity of each item in the package they desire. The proxy then bids in an ascending package auction.

In our application, the proxy phase of the auction consists of multiple rounds. The auction favors bids for long runs of consecutive slots when the service is provided by the same coalition. This strategy is designed to exploit temporal and spatial locality.

The auction starts with the longest runs and the lowest price per slot and proceeds with increasingly shorter runs and diminished incentives. Once a run of consecutive slots is the subject of a provisional winning bid, all shorter runs of slots for that particular service are removed from the coalition offerings.

During the first round only the longest run of consecutive slots for each one of the services offered by the participating coalitions is auctioned and only bidders that have committed to any of the slots of the run are allowed to bid. The price per slot for the entire run is the lowest price for any slot of the run the bidder has committed to during the clock phase of the auction. If there are multiple bids for service S_l^t the *provisional winner* is the one providing the largest revenue for the coalition offering the service.

If κ_l^t is the longest run of consecutive slots for service S_l^t auctioned in the first round then, in the second round, a shorter run of $\kappa_l^t - 1$ slots is auctioned. The price for the entire run equals the second lowest price for any slot of the run the bidder has committed to during the clock phase of the auction times the number of the time slots in the run.

The length of the consecutive slot runs auctioned decreases and the incentives diminish after each round. The preliminary rounds end with the auction of a single slot for each service. At the end of the preliminary round each bidder is required to offer the price for the slot committed to during the clock phase. Figure 5 depicts a plausible snapshot at the end of the preliminary rounds of the proxy phase when four services $\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3$ and $\mathcal{S}_4$, are offered and shows the provisional winners for service $\mathcal{S}_4$.



Figure 5: A snapshot at the end of the preliminary rounds of the proxy phase when there are four services offered and the auction covers 18 allocation slots. Dotted lines represent the quantity of service with provisional winners. Only the provisional winers for S_4 are shown, the clients labeled as $C_9, C_{13}, C_6, C_1, C_{11}, C_7$ and C_9.

During the final round the bidders reveal the packages they want to reserve; these packages include only the provisional winners from the preliminary slots. Once all provisional winning bids for services in a reservation request are known, the auctioneer chooses the package that best matches the consumer's needs and, at the same time maximizes the profit for the cloud service provider. The *coalition* for a reservation request consists of the set of coalitions that provide the services in the winning package.

In this auction all bids are firm, they cannot be withdrawn. The auction is monotonic, the length of runs of consecutive slots auctioned decreases continually; this guarantees that the auction eventually terminates. Linear pricing guarantees that the price of any package can be computed with ease.

The effectiveness of the protocol is captured by several metrics including:

-1. The *customer satisfaction index* - percentage of reservation requests fully or partially satisfied in each allocation slot given the total number of requests.

-2. The *service mismatch index* - percentage of services requested but not offered in each allocation slot given the total number of services in that slot.

-3. The *service success index* - percentage of services used in each allocation slot given all services offered in that slot.

-4. The *capacity allocation index* - percentage of the capacity offered but not auctioned in each allocation slot given the capacity offered in that slot.

-5. The *overbidding factor* - percentage of slots with a provisional winner that have not been included in any package given all slots offered at the beginning of the auction.

-6. The *temporal fragmentation index* - percentage of services successfully auctioned in non-consecutive slots given all services successfully auctioned.

-7. The *additional profit index* - percentage of additional profit of coalitions involved in the auction (the difference of the actual price obtained at the auction and the price demanded by the coalition) relative to the price demanded by the coalition.

Limitations and vulnerabilities. The protocol is fairly complex and has at least one vulnerability. A bidder may be the provisional winner of services in slots not included in its winning package; such services will remain unassigned during the current auction. A solution is to penalize*excess bidding activity* and charge the bidder a percentage of the costs for these services. Another alternative is to include, in a reservation request, a set of "substitute services" for a service S_i. Then, during the last round of the proxy phase, the auctioneer could try to match services having provisional winners with unsatisfied requests for services.

The capacity offered, but not auctioned in each slot is available for *spot allocation* thus, it has the potential to be used, rather then being wasted. The capacity of a coalition left uncommitted at the end of the auction $\mathbb{A}^t$ for AS_1^t, the first slot of the auction, is then available for *spot allocation* at a price equal to $p_{k,l}$, while the free capacity in slots starting with AS_2^t can be offered at the next auction if this auction takes place before the beginning of the slot. This capacity is measured by the *spot allocation opportunity index.*

6 Protocol Analysis and Evaluation

We report on the results of our simulation experiments to gain some insight into the proxy phase of the clock-proxy auction of the PC^2P protocol. The system we wish to evaluate requires the description of the environment in which the auction takes place, the reservation requests, and the services offered:

1. The environment elements: n - the number of coalitions offering services in this round; m - the number of clients; and κ - the number of slots auctioned.

2. The package j requested by client i: α_i^n - the number of services in the package; the slots desired by the service S_k, ordered by the length of the run of consecutive slots; $r_{k,j}$ - the intensity of service S_k in slot j; $p_{i,j}$ - the price per unit of service for slot j if client i was a provisional winner of that slot during the clock phase.

3. The service S_k provided by coalition $\mathbb{C}_k$ includes: γ_k - the largest run of consecutive slots for each offered service S_k; the profile of the service S_k - the slots offered ordered by the length of consecutive slots, when it is available; $q_{k,j}$ - the quantity of service S_k offered in slot j; and p_k - the price per unit of service offered by coalition $\mathbb{C}_k$.

For simplicity, we assume that a coalition offers one service only and the number of services is $\nu < n$. We also assume that all platforms have a maximum capacity of 100 vCPUs and that $q_{k,j}$, the quantity of service S_k offered for auction, and $r_{k,j}$, the quantity of S_k requested in slot j are the same for all the slots of an offered/requested run. The number of slots auctioned is fixed, $\kappa = 50$.

The range and the distribution of parameters for the protocol evaluation are chosen to represent typical cases. The parameters of the simulation are random variables with a uniform distribution:

-a. The number of coalitions and clients requesting reservations, n and m, respectively; the interval is $[200-250]$.

-b. The number of services offered and requested ν; the interval is $[10-20]$.

-c. The number of clients bidding for each service in a given slot; the interval is $[0-4]$.

-d. The capacity offered for auction for a service in a given slot; the interval is $[60-90]$ vCPUs.

-e. The services offered by a coalition; the interval is $[1-\nu]$.

-f. The number of consecutive slots a service is offered in; the interval is $[1-\kappa]$.

-g. The number of services in the package requested by a client; the interval is $[1-3]$.

-f. The number of consecutive slots of the services in the package requested by a client; the interval is $[1-\kappa]$.

We also randomly choose the slots when the client is the provisional winner. The evaluation process consists of the following steps:

A. Initialization.

B. Preliminary rounds. Carry out γ preliminary rounds with $\gamma = \max_k \gamma_k$.

- In the first preliminary round auction auction κ_1 slots of service S_1, κ_2 slots of service S_2, and so on.
- Identify the first slot of each run and the reservation request that best matches the offer.
- Identify the provisional winners if such matches exist and remove the corresponding runs from the set of available runs. A match exists if the run consists of the same number of slots or is one slot longer than requested and if the capacity offered is at least the one required by the reservation request. For services without a match, remove the last slot, add both the shorter run and the last slot to the list of available runs.
- Continue this process until only single slots are available.

C. Final round. In this round we:

- Identify the packages for each client and if multiple packages exist determine the one which best matches the request.
- Compute the cost for the winning package for each client.

Figures 6(a)-(e) show several performance metrics including the customer satisfaction index, the service mismatch index, the auction success ratio, the spot opportunity index, the temporal fragmentation index, and the capacity allocation index. The simulation covers 50 time slots.

The 5% confidence intervals for the mean of all performance metrics are computed for 25 batches each one of 200 realization of each random variable. The simulation times are 6.4 seconds for $2,000$ runs and 11.7 seconds for $5,000$ runs. The confidence intervals are rather tight; this indicates that the performance of the protocol is relatively stable for the range of parameters explored in this evaluation.

The auction success rate is high, typically above 80%. The initial low auction success rate is an artifact of the manner we conducted the simulation; we picked up randomly the service start up time. The spot allocation opportunity index is in turn correlated with the auction success rate and shows that a significant fraction of the capacity is available for spot allocation. This result is correlated with the one in Figure 6(f) which shows that

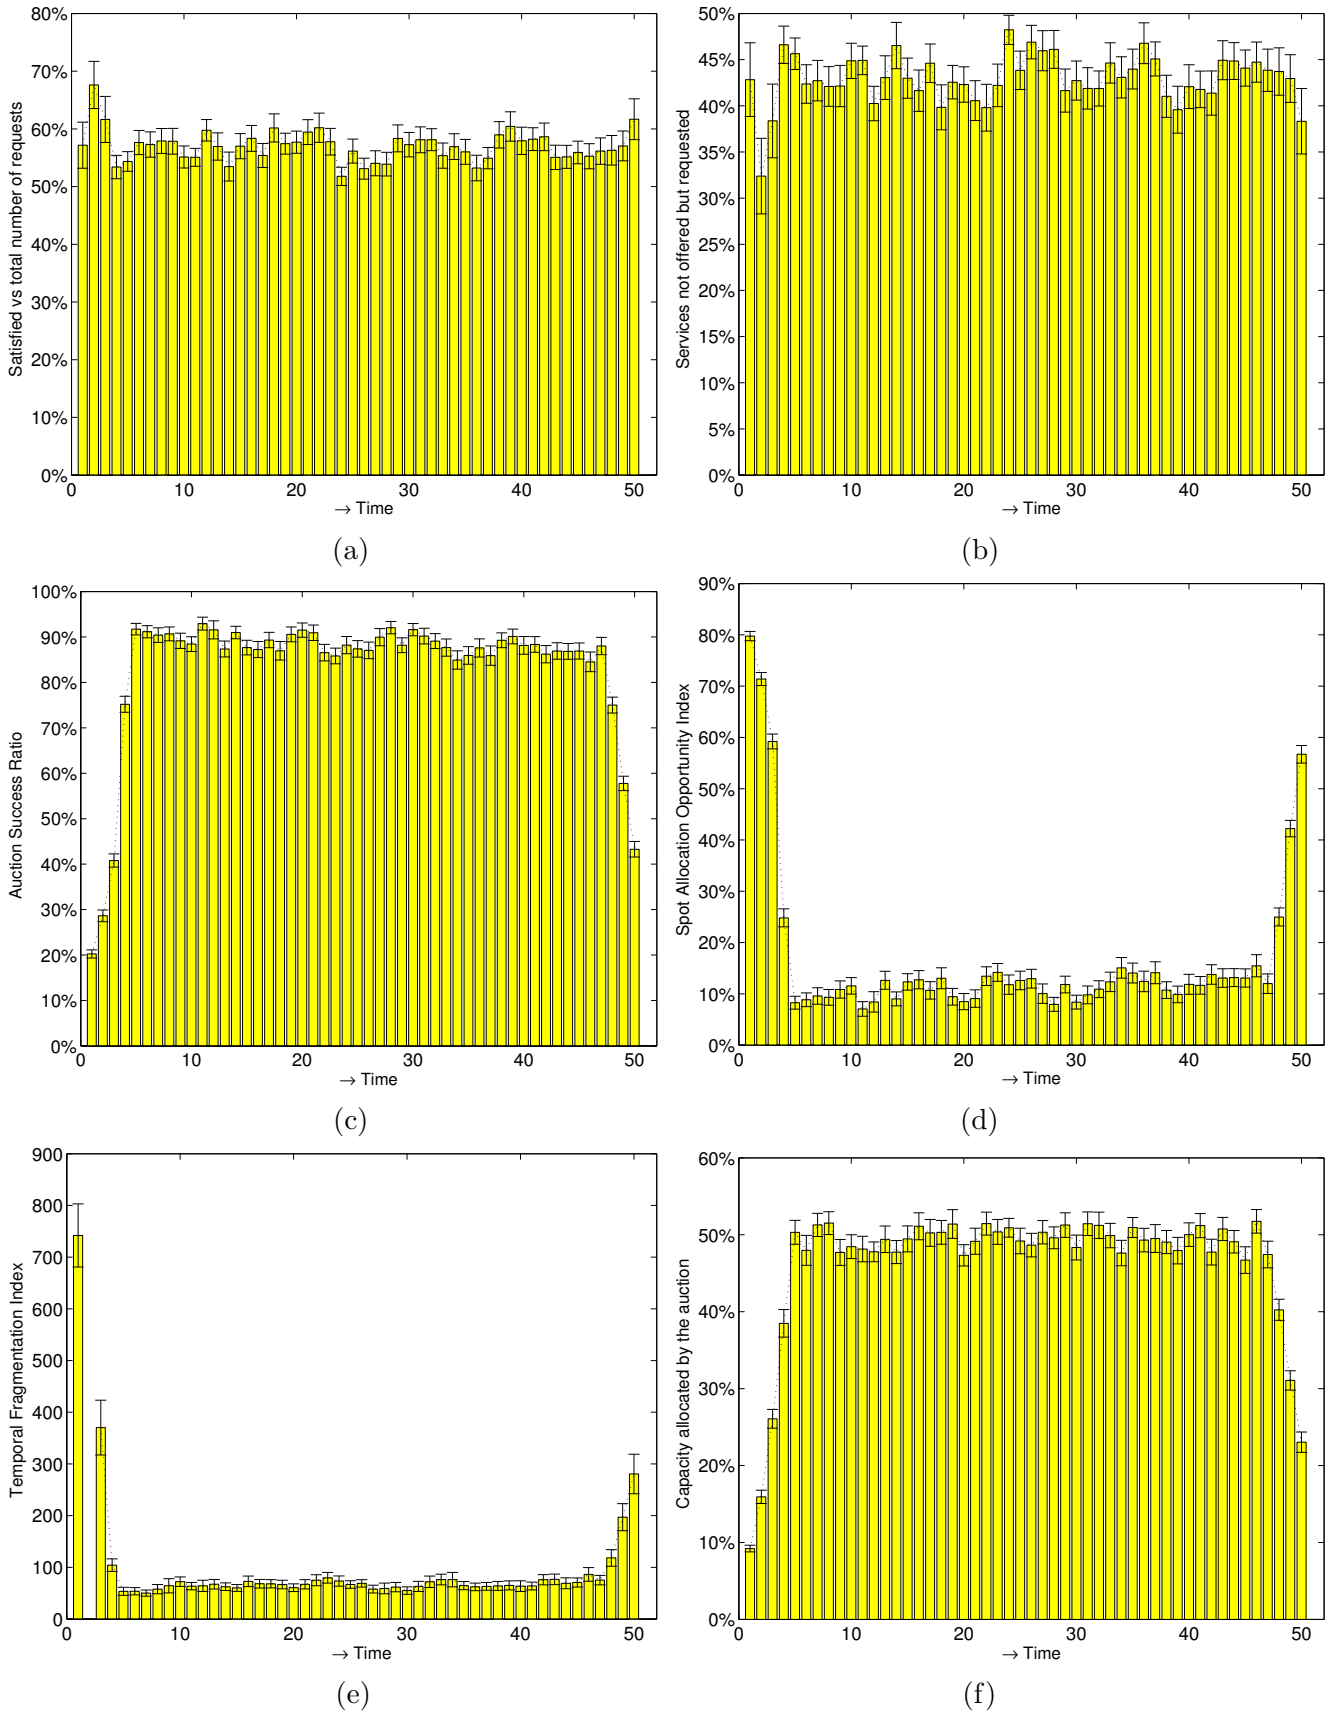


Figure 6: Proxy phase of an auction with 50 time slots. Indices of: (a) Customer satisfaction; (b) Service mismatch; (c) Auction success; (d) Spot allocation opportunity; (e) Temporal fragmentation; (f) Capacity allocation.

on average some 50% of the server capacity is not allocated by the reservation system and so is available for spot contention.

A reservation system covering 50% of the server capacity is probably the most significant result; it shows that self-management based on auctions can drastically improve server utilization. We live in a world of limited resources and cloud over-provisioning is not sustainable either economically or environmentally.

The service mismatch index is fairly high, typically in the 50% range and it is above 60% in a few slots. The customer satisfaction is correlated with the service mismatch and typically is in the region of 50%. In a realistic scenario, when coalitions maintain statistics regarding the services offered and avoid offering services unlikely to be demanded by the cloud users, the service mismatch would not affect the performance of the algorithm. Temporal fragmentation, though rather low, is undesirable. The overbidding factor $64 \pm 2.93\%$ is another indication that the protocol needs to be fine tuned.

Self-organization cannot occur instantaneously in an adaptive system and this simple observation has important consequences. It is critical to give autonomous cloud platforms, interconnected by a hierarchy of networks, the time to form coalitions in response to services demanded. Thus, self-management requires an effective reservation system and our results indicate that the reservation protocol is working well.

7 Conclusions and Future Work

Self-organization and self-management offer an appealing alternative to existing cloud resource management policies; they have the potential to significantly alter the cloud computing landscape. So far, pragmatic means for the adoption of self-organization principles for large-scale computing and communication systems have eluded us. A main reasons for this state of affairs is that self-management has to be coupled with some mechanisms for cooperation; these mechanisms should allow autonomous servers, to act in concert towards global system goals. Cooperation means that individual systems have to partially surrender their autonomy. Striking a balance between autonomy and cooperation is a challenging task, it requires a fresh look at the mechanics of self-organization and the practical means to achieve it.

Practical implementation of cloud self-organization is challenging for several reasons including the absence of a technically suitable definition of self-organization, a definition that could hint to practical design principles for self-organizing systems and quantitative evaluation of the results. Computer clouds exhibit the essential aspects of complexity and it is inherently difficult to control complex systems.

We started our investigation with a realistic model of the cloud infrastructure, the hierarchical organization reported in [5] which seems inherently tied to hierarchical control. First, we compared hierarchical control which based on monitoring with a market model in which the servers of a WSC place bids for service requests and found out that the latter is much more effective than hierarchical control [22]. In the simple market model the servers act individually, rather than cooperating with each other, a fundamental aspect of self-organization. But cooperation is clearly needed because individual servers may not be able to supply the resources demanded by many data-intensive application. Thus, we concluded that servers have to form coalitions to offer larger pools of resources. At the same time, it seemed obvious to us that complex applications with multiple phases would require packages of resources offered by different coalitions.

Algorithms for coalition formation based on combinatorial auctions are at the heart of the cloud ecosystem we propose. The path we chose seems logical as auctions have been successfully used for resource management in the past. Auctions do not require a model of the system, while traditional resource management strategies do. The auction-based protocol is scalable, and the computations can be done efficiently, though the computational algorithms involved are often fairly complex.

The results reported in Section 6 indicate that the performance of the protocol is relatively stable for the range of parameters explored in our evaluation. The protocol leads to a higher server utilization and it seems reasonable to expect that a fine-tuned version of the protocol could further improve this critical performance measure.

Two basic strategies for coalition formation are possible: (1) the one described in this paper, when coalitions are formed using information from previous combinatorial auctions, before knowing what the actual user demands are; (2) the more natural one, when coalitions are formed in response to user demands and then combinatorial auctions are organized. The former strategy, though simple and elegant, has two obvious drawbacks: (a) it needs the past history thus, starting from some arbitrary initial state may not work very well; (b) there may be cases when coalitions formed based on past history are not useful to any package and, at the same time, one or more packages could benefit from the available resources of the same type, but in coalitions of different sizes. The latter strategy, creating coalitions in response to known user requests is more complex. It requires a protocol to inform supply agents of the atrributes of desirable coalitions, a protocol for individual agents to express their willingness to join a particular coalition, and, finally, a coalition formation algorithm. Once the coalitions are formed the bidding for service packages can take place.

These two strategies can be combined, the latter used during the initial stages, when the historic information is either missing or incomplete, and the former used when history data is available. The second drawback of the strategy described in this paper can be attenuated by allowing a second round of coalitions formation. In this

second round the unsuccessful coalitions are disassembled and new coalitions matching the needs of unsatisfied packages are formed. For example, if two coalitions of size θ_1 and θ_2 were unsuccessful and a package $\mathcal{P}$ requireing a coalition of size $\theta_3 \leq \theta_1 + \theta_2$ was unsatisfied, then a second round will guarantee that $\mathcal{P}$ is satisfied and θ_3 servers will be reserved, instead of being left idele or offered to spot allocations.

Our future work will address this approach, as well as, other problems revealed by this investigation, e.g., the effects of overbidding. Overbidding is the process allowing a client to become a provisional winner of one or more service slots and then, in the final round failing to acquire some of them. This situation is critical for the first slot of an auction as the next auctions could find clients for these slots. A more difficult problem is the temporal fragmentation which does not seem to have an obvious solution.

References

[1] L. Ausubel and P. R. Milgrom. "Ascending auctions with package bidding." *Frontiers of Theoretical Economics*, **1**(1):1–42, 2002.

[2] L. M. Ausubel and P. Cramton. "Auctioning many divisible goods." *Journal European Economic Assoc.*, **2**(2-3):480-493, 2004.

[3] L. Ausubel, P. Cramton, and P. Milgrom. "The clock-proxy auction: a practical combinatorial auction design." *Chapter 5,* in *Combinatorial Auctions*, P. Cramton, Y. Shoham, and R. Steinberg, Eds. MIT Press, 2006.

[4] L. A. Barroso and U. Hözle. "The case for energy-proportional computing." *IEEE Computer,* **40**(12):33–37, 2007.

[5] L. A. Barossso, J. Clidaras, and U.Hözle. *The Datacenter as a Computer; an Introduction to the Design of Warehouse-Scale Machines.* (Second Edition). Morgan & Claypool, 2013.

[6] M. Blackburn and A. Hawkins. "Unused server survey results analysis." *www.thegreengrid.org/media/WhitePapers/Unused %20Server%20Study _WP_101910_v1. ashx?lang=en* (Accessed on December 6, 2013).

[7] D. Bruneo. "A stochastic model to investigate data center performance and QoS in IAAS cloud computing systems." *IEEE Trans. on Parallel and Distributed Systems,* **25**(3):560–569, 2014.

[8] T. E. Carroll and D. Grosu. "Formation of virtual organizations in grids: a game-theoretic approach." *Concurrency and Computation: Practice and Experience,* **22**(14):1972–1989, 2010.

[9] S. Chaisiri, B. Lee, and D. Niyato. "Optimization of resource provisioning cost in cloud computing." *IEEE Trans. on Services Computing,* **5**(2):164–177, 2012.

[10] V. Chang, G. Wills, and D. De Roure. "A review of cloud business models and sustainability." *Proc. IEEE 3rd Int. Conf. on Cloud Computing*, pp. 43–50, 2010.

[11] E. H. Clarke. "Multipart Pricing of Public Goods." *Public Choice*, **IX**:13–33, 1971.

[12] T. Groves. "Incentives in teams." *Econometrica,* **41**:617–631, 1973.

[13] J. O. Gutierrez-Garcia and K.- M. Sim. "Self-organizing agents for service composition in cloud computing." *Proc IEEE 2nd Int. Conf. on Cloud Computing Technology and Science,* pp. 59–66, 2010.

[14] L. He and T. R. Ioerger. "Forming resource-sharing coalitions: a distributed resource allocation mechanism for self-interested agents in computational grids." *Proc. ACM Symp. on Applied Computing,* pp. 84–91, 2005.

[15] S. U. Khan and I. Ahmad. "A cooperative game theoretical technique for joint optimization of energy consumption and response time in computational grids." *IEEE Trans. on Parallel and Distributed Systems,* **20**(3): 346-360, 2009.

[16] K. Lerman and O. Shehory. "Coalition formation for large-scale electronic markets." *Proc. ICMAS 2000 - 4th Int. Conf on Multiagent Systems,* pp. 167–174, 2000.

[17] C. Li and K. Sycara. "Algorithm for combinatorial coalition formation and payoff division in an electronic marketplace." *Proc. AAMAS02 - First Joint Int. Conf. on Autonomous Agents and Multiagent Systems,* pp. 120–127, 2002.

[18] H. Li, C. Wu, Z. Li, and F. Lau. "Profit-maximizing virtual machine trading in a federation of selfish clouds." *Proc. of the IEEE INFOCOM,* pp. 25–29, 2013.

[19] H C. Lim, S. Babu, J. S. Chase, and S. S. Parekh. "Automated control in cloud computing: challenges and opportunities." *Proc. First Workshop on Automated Control for Datacenters and Clouds,*, ACM Press, pp. 13–18, 2009.

[20] D. C. Marinescu, C. Yu, and G. M. Marinescu. "Scale-free, self-organizing very large sensor networks." Journal of Parallel and Distributed Computing (JPDC), **50**(5):612–622, 2010.

[21] D. C. Marinescu. *Cloud Computing; Theory and Practice.* Morgan Kaufmann, a division of Elsevier, Amsterdam, New York, 2013.

[22] D. C. Marinescu, A. Paya, J. P. Morrison, and P. Healy. "Distributed hierarchical control versus an economic model for cloud resource management." *http://arxiv.org/pdf/1503.01061v1.pdf*, March 2015.

[23] L.Mashayekhy, M.M. Nejad, and D. Grosu. "Cloud federations in the sky: formation game and mechanisms." *IEEE Trans. on Cloud Computing,* 2015 (to appear).

[24] I. Müller, R. Kowalczyk, and P. Braun. "Towards agent-based coalition formation for service composition." *Proc. IEEE/WIC/ACM Int. Conf. on Intelligent Agent Technology,* pp. 73-80, 2006.

[25] D.Niyato, A.Vasilakos, and Z.Kun. "Resource and revenue sharing with coalition formation of cloud providers: Game theoretic approach." *Proc. IEEE/ACM Intl. Symp. on Cluster, Cloud and Grid Comp.,* pp. 215–224, 2011.

[26] A. Paya and D. C. Marinescu. "Energy-aware load balancing and application scaling for the cloud ecosystem." *IEEE Trans. on Cloud Computing.* doi:10.1109/TCC.2015.2396059.

[27] S. Penmatsa and A. T. Chronopoulos. "Price-based user-optimal job allocation scheme for grid systems." *Proc of Parallel and Distributed Processing Symposium,* pp. 8-16, April 2006.

[28] T. Rahwan, S. D. Ramchurn, N. R. Jennings, and A. Giovannucci. "An anytime algorithm for optimal coalition structure generation." *Journal of Artificial Intelligence Research,* **34**:521–567, 2009.

[29] N. Samaan. "A novel economic sharing model in a federation of selfish cloud providers." *IEEE Trans. on Parallel and Distributed Systems,* **25**(1):12–21, 2014.

[30] T. W. Sandholm, K. S. Larson, M. Andersson, O. Shehory, and F. Tohm. "Coalition structure generation with worst case guarantees." *Artificial Intelligence,* **111**(1-2):209–238, 1999.

[31] S. Sen and P. S. Dutta. "Searching for optimal coalition structures." *Proc. ICMAS 2000 - 4th Int. Conf on Multiagent Systems,* pp. 287–295, 2000.

[32] M. Sims, C. V. Goldman, and V. Lesser. "Self-organization through bottom-up coalition formation." *Proc. Int. Conf. on Autonomous Agents and Multi Agent Systems,* pp. 867–874, 2003.

[33] R. Subrata, A. Y. Zomaya, and B. Landfeldt. "Game-theoretic approach for load balancing in computational grids." *IEEE Trans. on Parallel and Distributed Systems,* **19**(1):66–76, 2008.

[34] S de Vries and R. Vohra. "Combinatorial auctions; a survey." *INFORMS Journal of Computing*, **15**(3):284–309, 2003.

[35] H-J. Zhang, Q-H. Li, and Y-L. Ruan. "Resource co-allocation via agent-based coalition formation in computational grids." *Proc Second Int. Conf. on Machine Learning and Cybernetics,*, pp. 1936–1940, 2003.

[36] G. Wei, A. Vasilakos, Y. Zheng, and N. Xiong. "A game-theoretic method of fair resource allocation for cloud computing services." *The Journal of Supercomputing,* **54**(2):252–269, 2010.

Dan C. Marinescu. During the period 1984-2001 Dan Marinescu was an Associate and the Full Professor in the Computer Science Department at Purdue University in West Lafayette, Indiana. Since August 2001 he is a Provost Professor of Computer Science at University of Central Florida. He has published more than 220 papers in referred journals and conference proceedings and several books including: *Cloud Computing: Theory and Practice*, Morgan Kaufmann, 2013.

Ashkan Paya is a Ph.D. candidate in the Electrical Engineering and Computer Science Department at University of Central Florida pursuing his degree in Computer Science. He graduated from Sharif University of Technology in Tehran, Iran, with a BS Degree in the same major in 2011. His research interests are in the area of resource management in large-scale systems and cloud computing.

John Morrison is the founder and director of the Centre for Unified Computing. He is a co-founder and director of the Boole Centre for Research in Informatics, a principle investigator in the Irish Centre for Cloud Computing and Commerce and a co-founder and co-director of Grid-Ireland. Prof. Morrison has held a Science Foundation of Ireland Investigator award and has published widely in the field of Parallel Distributed and Grid Computing. He is a principle investigator in the Irish Centre from Cloud Computing and Commerce, where he leads the Service LifeCycle Group. He has been the guest editor on many journals including the Journal of Super Computing and the Journal of Scientific Computing. He is on the Editorial Board of Multi-Agent and Grid Systems: An International Journal, published by ISO Press, and the International Journal of Computational Intelligence: Theory and Practice (IJCITP). He is a senior member of the ACM and a senior member of the IEEE.